arXiv:1607.02769v1 [cs.CV] 10 Jul 2016

Annotation Methodologies for Vision and Language Dataset Creation

Gitit Kehat
Computer Science Department
Brandeis University
Waltham, MA. 02453 USA
gititkeh@brandeis.edu

James Pustejovsky
Computer Science Department
Brandeis University
Waltham, MA. 02453 USA
jamesp@brandeis.edu

Abstract

Annotated datasets are commonly used in the training and evaluation of tasks involving natural language and vision (image description generation, action recognition and visual question answering). However, many of the existing datasets reflect problems that emerge in the process of data selection and annotation. Here we point out some of the difficulties and problems one confronts when creating and validating annotated vision and language datasets.

1. Introduction

Recently, the use of Natural Language Processing (NLP) resources has become increasingly popular among the Computer Vision (CV) community, mostly thanks to the large-scale, easily accessible data from the web, and the growing popularity of online crowdsourcing platforms, such as Amazon Mechanical Turk (AMT). Typical vision and language tasks that utilize such resources are image description generation, action and affordance recognition, and visual question answering (VQA). This expanding intersection has also been leading to the definition of new vision and language tasks, some of them, in their language-only form, have been well studied in the NLP community, such as VQA. In addition, some works aim to utilize both language and vision to create richer multi-modal semantic spaces and vectors [8].

Most previous work analyzing vision and language datasets has dealt with the technical aspects of the collected datasets, rather than the data-gathering and annotation techniques used.[1] Work about the annotation process of images mainly focused on the speed, efficiency and cost aspects of the process [9]. As far as we know, we are the first to discuss the issue of the quality of *annotation content* within vision and language datasets. We present some of the major difficulties involving building and validating annotated vision and language resources, discuss their potential effects on results, and comment on combining NLP resources.

[1]The survey in [4] presented a large listing of datasets, analyzed by number of images and structure and more advanced criterions like syntactic richness of the vocabulary and the density of the captions.

A binary annotation task from [2] for the object "bicycle". The action "straddle" was marked as "ambiguous" due to annotator disagreement.

2. Vision and language annotation tasks

Annotation tasks can have the form of multiple choice questions (closed) or open response ones, or a combination of the two. The first case includes choosing "yes"/"no" for a given option (e.g. a pair of an action and an object), choosing all true options from a given list (e.g. given an object, pick its attributes or actions) and more. The Second case includes supplying a full free-form sentence describing an image, picking free words to describe attributes or actions, or fill-in-the-blank a specific attribute or event [15]. A combined case can allowed the user to add their own option to the list if needed [13]. Tasks can ask the user to refer to the whole image or to only certain parts or objects, and sometimes require them to annotate the regions of some relevant objects. The images can be natural scenes taken by real people, or artificial scenes created with clip arts like in [1].

2.1. Dimensions for Comparison

From the resources we surveyed, several dimensions emerged that can introduce potential weaknesses and inconsistencies during the creation of annotated datasets.

1. **Manual processing.** This concerns both creating and validating the gathered data. Manual gold sets are widely used in the verification process. In [2] for example, the set of verbs for each object is filtered manually, and verbs with similar meaning are grouped together manually. In [6], verbs for annotation are chosen and filtered manually, and annotations are examined manually too in order to penalize (what the authors see as) common mistakes. Apart from being inefficient and unscalable, this also creates author bias.
2. **Author bias.** This can occur when the gold set to evaluate annotation is created solely by the paper's authors. A

way to minimize the bias is to start with a small set and bootstrap it according to an annotation guideline, as in [6], or to use as a gold set the majority of a subset of the annotations, as in [5].

3. **Limited or sparse vocabulary.** In [2] and [6] annotators can respond only to options created by a predefined set of verbs/actions, and have no way of introducing new terms into the dataset, imposing a heavier burden on initial annotation schema design. A possible solution can be to let the users add their own action terms, if necessary, as in [13]. Not limiting the users at all, on the other hand, can potentially create a dataset with very low and insignificant counts for each term. This can become a problem when the dataset is small, like in [11]. Grouping together similar terms is one way to solve it. However, such unification can potentially lead to a lose in the fine differences between textually described actions. An immediate preferable solution would be to largely extend the size of the dataset.

4. **Action/visual sense is not well defined.** This becomes a problem when dealing with action or affordance recognition. A well-defined annotation is necessary for both grounding the action to match external evaluations and resources, and for creating consistency among annotators of the same dataset. This is especially crucial when dealing with data sets that were annotated with binary choices, such as [2, 6] (see Figure 1).

5. **Annotator's Attention.** When asked to describe an image, people tend to pick easy-to-describe relation (like "man wearing a t-shirt") and start with the most salient parts of the image[10]. In [11] and [13] workers were asked to annotate action for specific already detected objects, forcing them to focus on objects otherwise might have been forgotten.

6. **Validation and Averaging.** These post-processing steps are especially important when the number of annotators per image is small. Most of the work surveyed validated the annotations to avoid data that was corrupted for various reasons. In [7], a short quiz verified the English level, and in [6] two sets of test annotation questions were used to verify accuracy in submission time as well as to filter out malicious turkers. Sometimes an averaging step is performed to simplify data from closed annotation tasks. However, when the averaging results in picking the majority of the annotators only (like in [6]), or in a general non-informative "ambiguous" tag (as in [2]), or when non-agreement is completely ignored as in [12], this can lead to a possible information loss - just in the more interesting cases. A better solution would be to weigh according to several confidence levels chosen by the user for each annotation task (vs. just Y/N) (as in [3]), or according to the frequency (agreement) among annotators (as in [13]).

3. Using NLP Resources

Many of the more recent sources we reviewed used NLP tools and resources in some capacity. They are usually deployed as a "single application" solution, in order to expand the vocabulary or enrich the vision-similarity equations with semantic data, without further use. However, it is important to understand the expected behavior of such tools. For example, semantic vector space models (like Word2Vec) cannot distinguish between synonyms and antonyms, since they tend to appear in the same context [14]. This is particularly important when dealing with binary attributes and yes/no questions.[2]

Acknowledgements: This work was supported by Contract W911NF-15-C-0238 with the US Defense Advanced Research Projects Agency (DARPA) and the Army Research Office (ARO).

References

[1] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. Lawrence Zitnick, and D. Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2425–2433, 2015.

[2] Y.-W. Chao, Z. Wang, Y. He, J. Wang, and J. Deng. Hico: A benchmark for recognizing human-object interactions in images. In *The IEEE Int. Conf. on Computer Vision (ICCV)*, December 2015.

[3] Y.-W. Chao, Z. Wang, R. Mihalcea, and J. Deng. Mining semantic affordances of visual object categories. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4259–4267, 2015.

[4] F. Ferraro, N. Mostafazadeh, T.-H. K. Huang, L. Vanderwende, J. Devlin, M. Galley, and M. Mitchell. A survey of current datasets for vision and language research. In *Conference on Empirical Methods in Natural Language Processing*, 2015.

[5] S. Gella, M. Lapata, and F. Keller. Unsupervised visual sense disambiguation for verbs using multimodal embeddings. *arXiv preprint arXiv:1603.09188*, 2016.

[6] S. Gupta and J. Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015.

[7] M. Hodosh, P. Young, and J. Hockenmaier. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research*, pages 853–899, 2013.

[8] S. Kottur, R. Vedantam, J. M. Moura, and D. Parikh. Visual word2vec (vis-w2v): Learning visually grounded word embeddings using abstract scenes. *arXiv preprint arXiv:1511.07067*, 2015.

[9] R. Krishna, K. Hata, S. Chen, J. Kravitz, D. A. Shamma, L. Fei-Fei, and M. S. Bernstein. Embracing error to enable rapid crowdsourcing. *arXiv preprint arXiv:1602.04506*, 2016.

[10] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *arXiv preprint arXiv:1602.07332*, 2016.

[11] D.-T. Le, J. Uijlings, and R. Bernardi. Tuhoi: Trento universal human object interaction dataset. *V&L Net 2014*, page 17, 2014.

[12] H. Pirsiavash, C. Vondrick, and A. Torralba. Inferring the why in images. *arXiv preprint arXiv:1406.5472*, 2014.

[13] M. R. Ronchi and P. Perona. Describing common human visual actions in images. *arXiv preprint arXiv:1506.02203*, 2015.

[14] E. Santus, Q. Lu, A. Lenci, and C.-R. Huang. Taking antonymy mask off in vector space. In *Proc. of PACLIC*, pages 135–144, 2014.

[15] L. Yu, E. Park, A. C. Berg, and T. L. Berg. Visual madlibs: Fill in the blank description generation and question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2461–2469, 2015.

[2] The only work we found to have mentioned this issue was VQA [1].